

Advanced Programming Platform for efficient use of Data Parallel Hardware

Luis Cabellos

Institute of Physics of Cantabria (IFCA), CSIC-UC

Santander, 39005, Spain

Email: cabellos@ifca.unican.es

Abstract—Graphics processing units (GPU) had evolved from a specialized hardware capable to render high quality graphics in games to a commodity hardware for effective processing blocks of data in a parallel schema. This evolution is particularly interesting for scientific groups, which traditionally use mainly CPU as a work horse, and now can profit of the arrival of GPU hardware to HPC clusters. This new GPU hardware promises a boost in peak performance, but it is not trivial to use. In this article a programming platform designed to promote a direct use of this specialized hardware is presented. This platform includes a visual editor of parallel data flows and it is oriented to the execution in distributed clusters with GPUs. Examples of application in two characteristic problems, Fast Fourier Transform and Image Compression, are also shown.

I. INTRODUCTION

The game industry saw in the 2000s the revolution of the programmable shaders. Programmable shaders insert specific code in the 3D graphics pipeline in order to have customized effects¹. Initially programmable shaders allowed to change the pixels color, but soon they evolved to be able to modify also the geometry, and currently they have the flexibility to interact with almost all graphical elements, achieving the so called General-Purpose computing on graphics processing units (GPGPU).

The use of GPGPU allow developers to program general processing applications using graphics hardware. The game companies soon noticed the problem with the programmable shaders: the graphics are the objective of visual artists, but programming is a hard and time consuming process that is not taught in art academies. And programmers do not have the experience, nor the knowledge to get the best visual result programming customized effects with graphics hardware. The solution reached in the game industry was to develop powerful tools, easy to use for the artist but flexible enough to get all the results that a programmer can implement directly. These tools are based on a data-flow programming design that allows the visual artist to create special effects and display them on-screen during editing, exactly as they will appear along game execution. The tools include the possibility of directed edition showing immediately the effects of a proposed change, and also the intermediate state of data (as graphics primitives) in the workflow and so to understand the result of the change as a whole product or as a sum of transformations.

Scientific communities have increasing needs for processing large amounts of data as fast as possible [1]. Under this demand, they see GPGPU hardware with its high peak processing power as a valuable resource to handle for their workflows [2], but they are now in a similar situation to the the game industry with the programmable shaders: it is not trivial to program GPU resources for use in real applications. Efficient use of GPUs is difficult because although it is possible to use them in almost every kind of algorithms, only a few of them are executed more efficiently than in a CPU with a more general architecture. The input data needs to be sent from the CPU to the GPU and results returned back, but there is a limited memory bandwidth between both processors. Also the internal bus in a GPU is more powerful, but it has a different organization of memory caches and locations compared to a CPU, and programs need to take this organization into account to benefit. Also the GPU has the advantage of large parallelism thanks to hardware replication, but it has limited and strict pipelines limiting context switching between tasks without suffering performance degradation.

Data-flow programming is a paradigm that constructs applications as directed graphs [3] [4]. The vertex of the graphs are processes and the edges between vertexes define the input/output of such processes and the path the data should travel. The applications in this paradigm are defined changing the set of vertexes and creating the network of edges between them. Starting from a clever definition of vertexes and changing only the edges, it is possible to create many different applications under this paradigm.

The data-flow programming paradigm is close to the preferred work methodology used by scientific communities. They usually share the data from experimental sources and define scientific workflows to analyze that data. Scientific groups with better computing skills also share computational services to analyze data. However researchers do not usually apply the data-flow paradigm directly to program their applications, they rather compose them through successive filter and processing steps starting from the raw data obtained from the experimental setup.

However, there are several solutions to compose workflows in a scientific programming context, starting with the popular and powerful LabView software [5], the Kepler System [6], or others adapted to a Grid computational infrastructure [7].

The proposal presented here aims to use the data-flow

¹Effects not available in the graphics hardware.

paradigm starting at the basic level, constructing modules from a well defined set of processes, conceived as orthogonal components, including data parallelism, and that communicate between them to build applications.

This Distributed Programming Platform for Data Parallel Algorithms will also benefit of powerful visual tools like those already developed under the data-flow paradigm [8]. These tools will include a user friendly editor to program data-flow applications when the algorithm fits into a data parallelism model, and also a service able to execute the resultant program flows in the most efficient way using computers with one or more GPUs.

The architecture explained in this paper is designed to allow the user to build the flow once and be able to execute it with different data sets and on different distributed computing hardware offering GPU resources.

II. IMPLEMENTATION

A. GPU Framework

The first step in the development of a Data-Parallel Platform is the selection of a GPGPU platform. Currently there are two major platforms available: CUDA and OpenCL [9]. CUDA stands for Compute Unified Device Architecture, and it is a computing engine developed by Nvidia Corporation enabling access to Nvidia GPUs as a GPGPU platform [10]. OpenCL is an specification made by the Khronos Group, of an open standard for general purpose parallel programming [11]. Although both platforms are presented to use GPU hardware as GPGPU platforms, they also support manycore and multicore hardware, so they are able to execute code in both CPUs and GPUs, provided the corresponding driver.

CUDA offers higher quality libraries and also includes better development tools, like a debugger and an emulator. Applications in CUDA require less setup code, and the performance compares favorably with OpenCL [12].

Both platforms use the C language as their reference model, and have similar memory and concurrency characteristics, so converting programs between both platforms is not difficult.

On the other hand, OpenCL has a clear advantage over CUDA: while CUDA is designed to work with NVIDIA hardware, OpenCL, as an open standard, has already drivers implemented for NVIDIA, ATI and Intel hardware, both for GPUs and CPUs.

From the point of view of a potential user of the Data-Parallel Platform, this is the most relevant argument. In fact, potential users of the Data-Parallel-Platform are not interested in the tools that the developers will use, nor in the setup details, but are critically concerned about availability of the platform for their existing hardware. Based on these arguments, OpenCL was selected as the platform to build the Data-Parallel Platform.

The Data-Parallel Platform uses OpenCL in two different ways. Firstly, the OpenCL software development kit is used to execute the program flows in the GPU hardware, including manycore and multicore hardware when available

<pre>for(int i = 0 ; i < MAX ; i++){ z[i]=x[i]+y[i]; }</pre>
<pre>__kernel adder(global float * x, global float * y, global float * z){ int i = get_global_id(0); z[i]=x[i]+y[i]; }</pre>
<pre>"adder":{ "body": "int i = get_global_id(0);\nz[i]=x[i]+y[i];\n", "io":{ "x":{ "data":"float", "type":"InputPoint"}, "y":{ "data":"float", "type":"InputPoint"}, "z":{ "data":"float", "type":"OutputPoint"}} }</pre>

TABLE I
COMPARATIVE OF THE IMPLEMENTATION OF A LOOP IN THE DIFFERENT PLATFORMS. TOP: THE LOOP IN C-LIKE PSEUDOCODE. MIDDLE: THE LOOP USING AN OPENCL KERNEL. BOTTOM: THE LOOP IN DATA-PARALLEL JSON FORMAT.

Secondly, the OpenCL C programming language is employed to codify the behavior inside the vertexes of the Data-Parallel Program graphs. As it will be shown below, there is a direct translation between the Data-Parallel Platform vertexes and OpenCL C source code. A simple example can be seen comparing middle and bottom sections in Table I.

This strategy can be seen as a complexity reduction of the access to GPU programming [13]: hand coding complexity for the final user is much reduced as OpenCL functions input and output parameters are limited. However it implies also a limitation on the complexity of the problem being coded, and addressed.

OpenCL support two execution models, data parallel and task parallel programming models. The Data-Parallel Platform will use only the data parallel programming model of OpenCL to offer data-flow programming to users. In this data parallel model, a sequence of instructions is applied to multiple elements in memory. Each one of these elements is called a work-item and the parallelism is achieved executing the sequence of instructions at the same time over all the work-items. In the Data-Parallel Platform a one-to-one bind between the work-item in memory and the kernel currently executed is established. In this way the input data-flow in a Data-Parallel program is split into chunks of work-items, then executed in parallel using OpenCL and finally the result is re-joined to compose the output data-flow.

B. The Data Parallel Model

The use of the data parallel model in OpenCL and the division in blocks of work-items requires that the Data-Parallel Programs are strictly Directed Acyclic Graphs (DAGs). This requirement avoids return edges, that would complicate the parallelism of blocks of elements if a vertex would have to

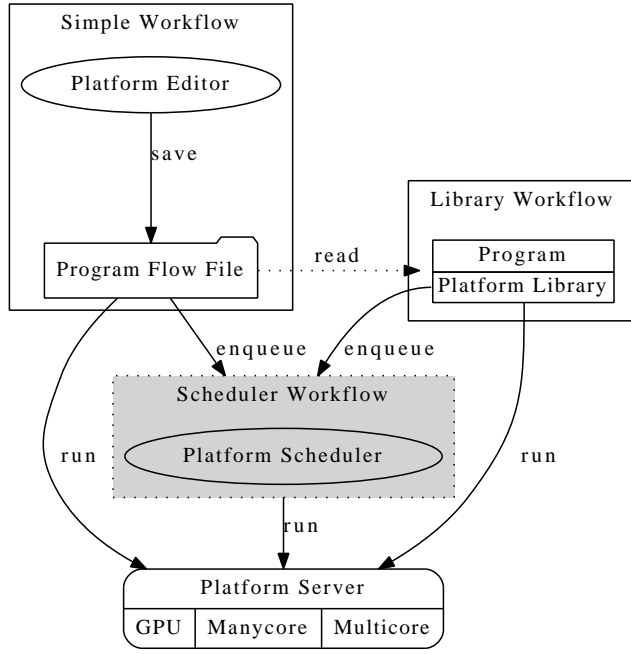


Fig. 1. Data-Parallel Platform workflow scheme. The graph shows the two different ways to use the Data-Parallel Platform: 1) As a library from a user application; 2) from the Data-Parallel Platform Editor, either running directly in the Data-Parallel Platform Server, or, in the future, as a job system.

wait for the output of a posterior vertex.

The Data-Parallel Platform is structured as several layers or components in order to allow the execution of Data-Parallel programs in different ways: direct execution, scheduled execution on a queue, or integrated in existing applications using a library. These different possibilities are represented in figure 1.

The most basic example would start with the creation of a Data-Parallel Program using the Data-Parallel Editor (see below), then selecting the input files to be processed, and finally executing the program in a Data-Parallel Server (also presented later).

The same Data-Parallel Program created using the editor can be executed by a program using the functions from the Data-Parallel Platform library. Execution using a queue system is being implemented under a wider scope, in a Distributed Data-Parallel Platform including a Data-Parallel Scheduler acting as a batch system for Data-Parallel Programs².

C. A Visual Editor of Data-Parallel Programs

A Data-Parallel Program Editor has been implemented as a visual tool following the Blender³ Compositor style. The edition of a complete Data-Parallel Program has two parts. The first one is the definition of the nodes: individual nodes are created and can be modified individually, including its input/output set and its body main program in OpenCL C Programming Language. The second part is the definition of

²The whole set of tools described here is being further developed under the name of the Skema Platform.

³Blender is the free open source 3D content creation suite.

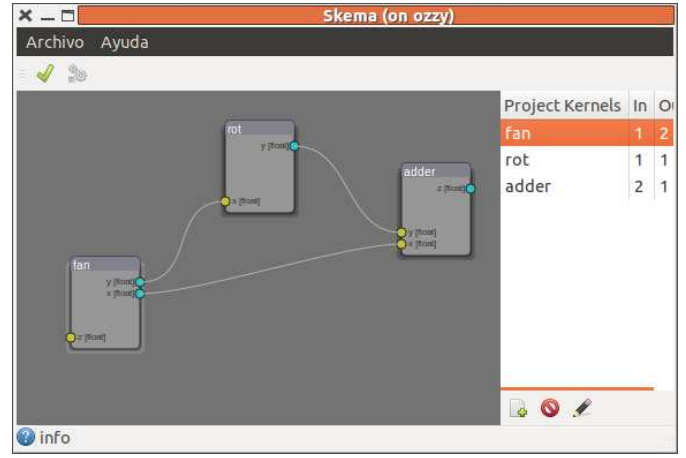


Fig. 2. The Visual Editor showing a basic Data-Parallel Program. The data flows from left to right with a floating number as input and a floating number as result.

the data flow between instances of the nodes. Nodes must be instantiated and arranged into a processing network.

Figure 2 shows how a Data-Parallel Program appears in the visual editor.

A Data-Parallel Program is a data-flow application created to be executed in the Data-Parallel Platform. Its main components are the following ones:

Type The available data types in the Data-Parallel Platform: OpenCL 1.0 [11] data types are used, including scalar and vector data types.

Input/Output Point Points attached to vertexes in the Data-Parallel Programs. The set of points of a Node define the possible communication channels between instances of that node.

Node A node defines the behavior of the graph vertexes in a Data-Parallel Program. It is composed of a set of Input/Output points (at least one of each type) and a main program body coded using the OpenCL C Programming Language specification.

Instance An instance of a node is a vertex in the Data-Parallel Program. In the example shown in the previous figure 2 there are three instances of three different nodes.

Arrow An arrow is an edge between two instances or vertexes. Specifically, an arrow connects an output point of an instance with a compatible input point from a different instance. The points are compatible if they have the same base scalar type⁴. A point from an instance without a connected arrow is named an unassigned point or a free point.

Program A Program is the directed acyclic graph of instances and arrows than can be executed in the Data-Parallel Platform. The diagram shown in figure 2 shows a basic Data-Parallel Program.

Stream A stream is a continuous flow of data with a de-

⁴The base type of a scalar is that scalar data type. The base type of a vector data type is the scalar element of the vector

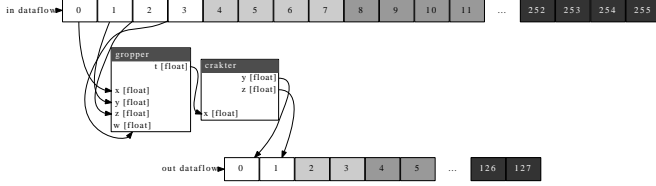


Fig. 3. Representation of the execution of a Data-Parallel Program. The Data-Parallel Program get chunks of data from an input stream, executes the programming code included in the nodes in parallel for each of the elements of that chunk, and generates an output stream composed of the results re-joined in adequate order.

defined type and related to a free point of a Data-Parallel Program. The execution of a program requires one or more input streams and one or more output streams. The figure 3 describes the execution of a Data-Parallel program over an input stream to generate an output stream.

Data-flow The Data-flow is the set including the input stream and the output stream in a Data-Parallel Program execution.

Continuing with the editing process, once the nodes, instances and arrows between instances of a Data-Parallel Program are defined, all the corresponding information is exported to a JSON [14] file. This file stores this information in a format used in the execution of Data-Parallel programs, either for sending it to a Data-Parallel Server or for connecting to the Data-Parallel Platform Library. The Table II shows a basic example of exported information corresponding to the Data-Parallel Program described in Figure 2. The corresponding JSON file is used by the Data-Parallel Platform to execute the program flow.

D. The Data-Parallel Server

The Data-Parallel Server is the module in the platform that executes the Data-Parallel programs on an input data-flow to obtain an output data-flow. For that reason it is the only module that actually requires the OpenCL driver and also direct access to the associated hardware. The server is in charge of communicating the state of the OpenCL platform, the state of the GPGPU hardware and its characteristics, and also the running progress of Data-Parallel programs. It uses a simplified approach for external communication based on REST (Representational State Transfer) [15] with HTTP networking protocol and JSON documents as information exchange format. The Data-Parallel Server currently is not RESTful as not all the REST architectural elements are implemented, in particular the layered and cacheable properties, although it is planned to include these features when possible in future developments to improve its scalability.

The Data-Parallel Server executes the Data-Parallel Programs using a simple Run Protocol to connect with the clients. As shown in Figure 4, this protocol defines the order of the steps to execute a program over a data-flow: send the Data-Parallel Program to the server, initialize the execution of the

```

"kernel": {
  "body": "int i=get_global_id(0);
  z[i]=x[i]+y[i];",
  "io": {
    "x": { "data": "float", "type": "InputPoint" },
    "y": { "data": "float", "type": "InputPoint" },
    "z": { "data": "float", "type": "OutputPoint" }
  }
},
"fan": {
  "body": "int i=get_global_id(0);
  x[i]=z[i].x;
  y[i]=z[i].y;",
  "io": {
    "x": { "data": "float", "type": "OutputPoint" },
    "y": { "data": "float", "type": "OutputPoint" },
    "z": { "data": "float2", "type": "InputPoint" }
  }
},
"rot": {
  "body": "int i=get_global_id(0); ny[i]=x[i]<<16;",
  "io": {
    "x": { "data": "float", "type": "InputPoint" },
    "y": { "data": "float", "type": "OutputPoint" }
  }
},
"nodes": [
  [0, { "kernel": "fan" }],
  [1, { "kernel": "rot" }],
  [2, { "kernel": "add" }],
],
"arrows": [
  { "output": [0, "x"], "input": [2, "x"] },
  { "output": [1, "y"], "input": [2, "y"] },
  { "output": [0, "y"], "input": [1, "x"] }
]

```

TABLE II

JSON FORMAT CORRESPONDING TO A DATA-PARALLEL PROGRAM. THIS BASIC EXAMPLE CORRESPONDS TO THE PROGRAM PREVIOUSLY SHOWN IN FIGURE 2, COMPOSED OF THREE NODES AND THREE INSTANCES, AND THE CORRESPONDING DATA-FLOW BETWEEN INSTANCES COMPOSED OF THREE EDGES.

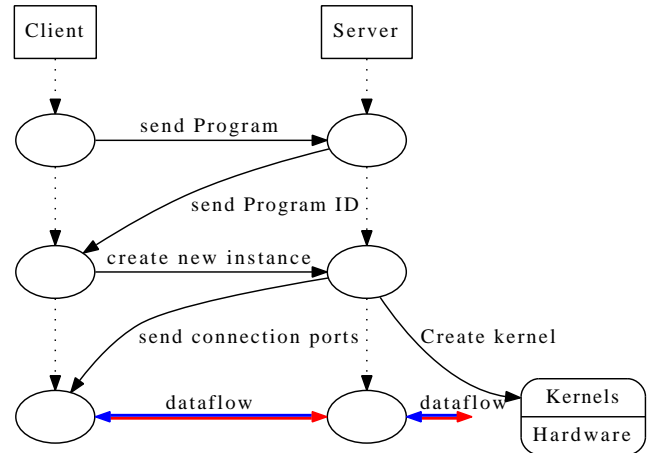


Fig. 4. Distributed Run Protocol. Client starts a running instance in the server, to execute the desired program, using a web API in http with JSON as data language; once the instance is created, the client sends and receives data from the server via the TCP protocol.

program and finally send the input data-flow and receive back the output data-flow. It is important to notice that the first step could be skipped if the program is transferred previously. For this purpose, an unique ID can be associated with the JSON representation of the program, a program ID. This option may save a significant time if the same Data-Parallel Program is to be executed with different input streams.

III. EXAMPLES

In order to show the capabilities of our Data-Parallel Platform two working examples are presented below. The first example computes the discrete Fourier transform using the Cooley-Tukey algorithm. The second example is a simple lossy image compression application using a visual vector quantization of blocks.

Both examples were tested in a server node running the Data-Parallel Server and a desktop computer running the client programs. The server node was a Megware Computer solution equipped with an Intel Xeon X5550 2.66GHz Quad Core processor and 4GB memory. It had four NVIDIA Tesla C1060 GPUs installed.

The desktop computer was an HP Proliant ML330 G6. Both computers are interconnected using a Gigabit Ethernet LAN.

A. Discrete Fourier Transform Example

The discrete Fourier transform (DFT) is a mathematical transform of a signal between discrete domains used for Fourier analysis. The DFT is widely used in signal processing, to analyze the frequencies of a signal, data compression eliminating frequencies with less information in a signal, polynomial multiplication and convolutions. All these applications depend upon an efficient calculation of this transformation, so it is a good example to test the speedup of this DFT using the Data-Parallel Platform on GPU hardware.

The Cooley-Tukey [16] algorithm was used to implement the Fast Fourier Transform. This algorithm, one of the most used in DFT, recursively calculates a DFT of N elements using two DFT of sizes N_1 and N_2 having $N = N_1 \cdot N_2$. When N is highly composite ⁵ the DFT computation time can be reduced from $O(N^2)$ to $O(N \cdot \log N)$. In particular, the radix-2 Cooley-Tukey algorithm was used, where the decimation of the DFT is done with two interleaved DFT of size $N/2$ in each recursive step.

Using this radix-2 decimation, the computation of the last k steps can be sent to the Data-Parallel Platform Server, parallelizing the calculus of a large number of DFT of size 2^k . The 2^k DFT is computed in a simple Program Node and the DFT flow is the input flow of the Data-Parallel Program.

Figure 5 presents the result of executing the Data-Parallel Program with a flow of DFT with sizes 2, 4 and 8 and the execution of the same data with a CPU implementation of the Cooley-Tukey algorithm for the same sizes. The time increase is linear in both implementations, as it should be, but it can be seen that the Data-Parallel Program is approximately five times faster, although it has to send the data over the local network to the Data-Parallel Server.

The Data-Parallel Program has also a remarkable advantage over the full CPU implementation. While the Data Parallel platform is executing the DFT calculations, the CPU usage in the local computer is only around 10% and corresponding to I/O time; in contrast the CPU implementation requires $\sim 90\%$ CPU usage.

⁵Highly composite numbers are numbers which factors completely into small prime numbers.

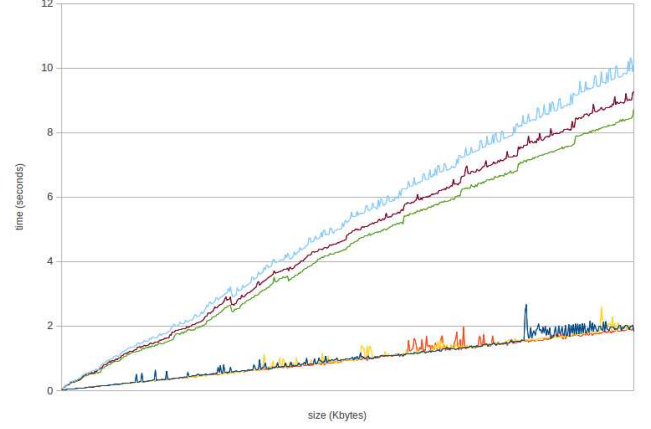


Fig. 5. Comparing DFT on CPU versus GPU execution. The size of the data used span from 20 Kbytes to 10 Mbytes. Three FFT sizes 2, 4 and 8 values in both types of tests (CPU and GPU) are used. All GPU tests remain well below 2 seconds in execution time while the CPU tests increase up to 8 \sim 10 seconds.

B. A second example: Image Block Compression

The purpose of image compression is to represent images using less data in order to save storage costs or transmission time. A raw image, without compression, can be quite large, usually up to several megabytes, and compression can reduce the file size very significantly. The image data can be compressed in such way that the exact original data can be recovered from the compressed data, or losing information in the image data but reaching better compression rates.

Lossy compression is usually based on techniques that remove details that humans do not notice. In this example a lossy image compression has been implemented using the Data-Parallel platform and employing well known methods. The first method used is the conversion of red-green-blue image data to a chrome-luminance representation, followed by color sub-sampling to scale 1/4 the from the full size using the fact that the human eye is more sensitive towards light intensity variation than color variation ⁶. The second method applied, explained in references [17] [18], divides the image luminance in blocks of 4×4 pixels, and determines a code book of N representative mean blocks. With this code book image blocks are encoded, using intensity deviation, instead of using the full information of the 16 pixels.

The implemented algorithm uses the following five steps:

- 1) Convert to Chrome + Luminance representation
- 2) Downscale Chrome layer
- 3) Calculate Directional Derivative of Luminance
- 4) Apply k-means for calculate codebook
- 5) Compress Luminance in blocks

Steps 1, 2 and 3 are calculated in the Data-Parallel Platform, while the creation of the code book (step 4) is made in the CPU with the returned data. Once the code book is created,

⁶Nobody will notice the colour downscale

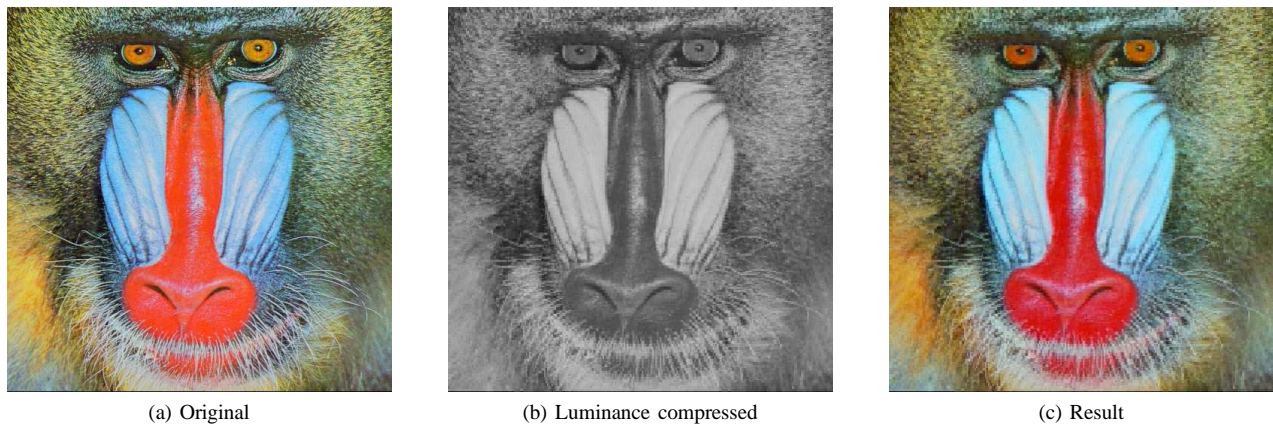


Fig. 6. Image compression from (a) original image to (c) compressed result. The block compression is done in the (b) luminance of the image, the colour layer is scaled down 1/4th from the full picture. Uncompressed size is ~ 770 Kbytes, and compressed size is ~ 80 Kbytes

the information is sent back to the GPU to calculate the compressed data.

The compression is noticeable, as can be seen in Figure 6, both in quality and in the reduction factor, but the objective with this example is to show how the Data-Parallel Platform can be used to build a working application. The sections of the program executed in the GPU were created using the visual editor, and during the execution of the image compression this work was distributed on a running Data-Parallel Server using the Data-Parallel Library.

IV. CONCLUSIONS AND OUTLOOK

A Data-Parallel Platform has been designed supporting the use of GPGPU on clusters allowing to access to the power of GPUs as a service, with the advantages this means for the implementation of work-flows and schedulers. Difficulties of GPGPU programming are reduced thanks to a clear programming model using the OpenCL platform and modeling the problems using DAGs, and offering a visual editor tool for final users powerful enough to exploit the GPGPU syntax.

There is an increasing trend to use GPUs specialized processors as common building blocks of supercomputers. China's Tianhe-1A supercomputer achieved in October 2010 the number one in the TOP500 ranking using graphics chips, and in the march 2011 3 of the top 5 supercomputers [19] were using mixed architectures with both CPUs and GPUs. Although the increase in performance thanks to the use of GPUs seems very high, with up to a 20x factor, GPUs require specialized programming, and the lack of advanced programming tools and languages with limited features is a problem [20].

The Data-Parallel Platform presented does not aim to be the best tool for performance, and it's not yet fully completed to offer all the characteristics planned, but it is a solution prepared to allow distributed computing with GPGPU hardware.

Many improvements will be required to make it a production tool. For example, regarding performance of program executions, the gap when using a cascade of instances due to inefficient movement of data between them, has to be solved

Graph theory must be revisited in order to further optimize the Data-Parallel Programs. In particular to understand how to split a Data-Parallel Program into several concurrent flows. There is also the possibility of include characteristics of other distributed solutions in the Data-Parallel Platform, like high availability, large scalability or Map/Reduce technologies.

Also the design of a job system for Data-Parallel Programs running on Data-Parallel servers, as part of a Distributed Data-Parallel Platform, will allow a better scale of applications and a better use of GPGPU resources, especially in computer clusters with GPU hardware.

ACKNOWLEDGMENT

This works was supported by the Ministry of Science and Innovation of Spain and their National Scientific Research, Development and Technological Innovation Plan (National R&D&i Plan) at the University of Cantabria.

REFERENCES

- [1] T. Hey and A. Trefethen, *The Data Deluge: An e-Science Perspective*. John Wiley & Sons, Ltd, 2003, pp. 809–824. [Online]. Available: <http://dx.doi.org/10.1002/0470867167.ch36>
- [2] Z. Fan, F. Qiu, A. Kaufman, and S. Yoakum-Stover, "Gpu cluster for high performance computing," in *Proceedings of the 2004 ACM/IEEE conference on Supercomputing*, ser. SC '04. Washington, DC, USA: IEEE Computer Society, 2004, pp. 47–. [Online]. Available: <http://dx.doi.org/10.1109/SC.2004.26>
- [3] Arvind and D. E. Culler, *Dataflow architectures*. Palo Alto, CA, USA: Annual Reviews Inc., 1986, pp. 225–253. [Online]. Available: <http://dl.acm.org/citation.cfm?id=17814.17824>
- [4] A. L. Davis and R. M. Keller, "Data flow program graphs," *IEEE Computer*, vol. 15, no. 2, pp. 26–41, 1982.
- [5] G. W. Johnson, *LabVIEW Graphical Programming: Practical Applications in Instrumentation and Control*, 2nd ed. McGraw-Hill School Education Group, 1997.
- [6] B. Ludscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. A. Lee, J. Tao, and Y. Zhao, "Scientific workflow management and the kepler system," in *Concurr. Comput. : Pract. Exper.*, 2005, p. 2006.
- [7] I. Foster, "What is the Grid? - a three point checklist," *GRIDtoday*, vol. 1, no. 6, Jul. 2002. [Online]. Available: <http://www-fp.mcs.anl.gov/~ifoster/Articles/WhatIsTheGrid.pdf>
- [8] W. M. Johnston, J. R. P. Hanna, and R. J. Millar, "Advances in dataflow programming languages," *ACM Comput. Surv.*, vol. 36, pp. 1–34, March 2004. [Online]. Available: <http://doi.acm.org/10.1145/1013208.1013209>

- [9] M. Harris and D. Gddecke. General-purpose computation on graphics hardware. [Online]. Available: <http://gpgpu.org>
- [10] *NVIDIA CUDA Compute Unified Device Architecture: Programming Guide*, 2nd ed., Nvidia Corporation, July 2008.
- [11] *The OpenCL Specification, Version 1.0*, Rev. 33 ed., Khronos Group Std., April 2009. [Online]. Available: <http://www.khronos.org/registry/cl/specs/opencl-1.0.33.pdf>
- [12] K. Karimi, N. G. Dickson, and F. Hamze, "A performance comparison of cuda and opencl," *CoRR*, vol. abs/1005.2581, 2010.
- [13] S.-Z. Ueng, M. Lathara, S. S. Baghsorkhi, and W.-M. W. Hwu, "Cuda-lite: Reducing gpu programming complexity," in *Languages and Compilers for Parallel Computing*, J. N. Amaral, Ed. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 1–15.
- [14] D. Crockford, "The application/json Media Type for JavaScript Object Notation (JSON)," RFC 4627 (Informational), Internet Engineering Task Force, Jul. 2006. [Online]. Available: <http://www.ietf.org/rfc/rfc4627.txt>
- [15] R. T. Fielding and R. N. Taylor, "Principled design of the modern web architecture," *ACM Trans. Internet Technol.*, vol. 2, pp. 115–150, May 2002. [Online]. Available: <http://doi.acm.org/10.1145/514183.514185>
- [16] J. Cooley and J. Tukey, "An algorithm for the machine calculation of complex fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965.
- [17] N. M. Nasrabadi and R. A. King, "Image coding using vector quantization: a review," *Communications, IEEE Transactions on*, vol. 36, no. 8, pp. 957–971, 1988. [Online]. Available: <http://dx.doi.org/10.1109/26.3776>
- [18] G. Qiu, "A fast algorithm for constructing image identification."
- [19] H. Meuer, E. Strohmaier, J. Dongarra, and H. Simon, "Top500 supercomputer sites list," TOP500.Org, Tech. Rep., June 2001. [Online]. Available: <http://top500.org/lists/2011/06>
- [20] P. Varhol, "Gpu vs. cpu computing: When your time is on the line, you need both types of processors." *Desktop Engineering*, September 2010. [Online]. Available: <http://www.deskeng.com/articles/aaayet.htm>